arXiv:1411.0802v1 [cs.RO] 4 Nov 2014

Simultaneous Localization, Mapping, and Manipulation for Unsupervised Object Discovery

Lu Ma, Mahsa Ghafarianzadeh, Dave Coleman, Nikolaus Correll, and Gabe Sibley*

***Abstract*— We present an unsupervised framework for simultaneous appearance-based object discovery, detection, tracking and reconstruction using RGBD cameras and a robot manipulator. The system performs dense 3D simultaneous localization and mapping concurrently with unsupervised object discovery. Putative objects that are spatially and visually coherent are manipulated by the robot to gain additional motion-cues. The robot uses appearance alone, followed by structure and motion cues, to jointly discover, verify, learn and improve models of objects. Induced motion segmentation reinforces learned models which are represented implicitly as 2D and 3D level sets to capture both shape and appearance. We compare three different approaches for appearance-based object discovery and find that a novel form of spatio-temporal super-pixels gives the highest quality candidate object models in terms of precision and recall. Live experiments with a Baxter robot demonstrate a holistic pipeline capable of automatic discovery, verification, detection, tracking and reconstruction of unknown objects.**

I. INTRODUCTION

Automatic object model acquisition is a fundamental task in robotics. Recent state-of-the-art object detection and recognition systems are based on pre-trained detectors [1]–[4] and features [5], [6]. For object learning, most techniques reconstruct sparse 2D models [7], [8] together with dense 3D shape models. The focus of this paper is an unsupervised pipeline for discovering unknown objects in a scene without any prior information and using robots with manipulators to verify and reinforce model acquisition by robot-object interaction (e.g. poking, grasping). Hence, a pre-trained object detection technique is not suitable here – indeed, the proposed method could be used to train such systems.

2D object appearance is highly informative for image based spatio-temporal segmentation [9], [10]. Typically, per-frame image segmentation results are combined with temporal information via 2D segment tracking. Instead, we find that using appearance and dense spatio-temporal cues produces better image segmentation results, which in turn yield a high quality set of likely candidate objects. However, these cues are not sufficient for the verification of object-hood, especially for static objects [11], [12]. In the proposed approach, a mobile robot relies on dense 3D SLAM to manipulate putative objects, thereby adding possible motion cues to verify object-hood.

Once an object has been verified by poking or grasping, a dense detection, tracking and reconstruction technique is applied to learn 2D and 3D implicit models of appearance and shape. In subsequent steps, learned models are automatically detected, tracked, and refined.

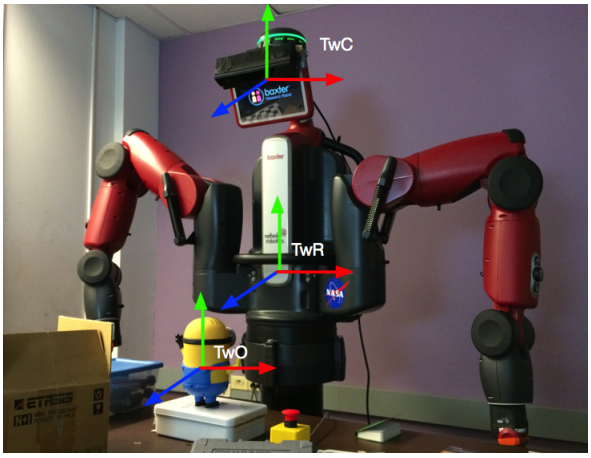

Fig. 1: The dual-arm Baxter robot equipped with a RGB-D camera automatically discovers and learns dense 3D model of unknown objects in the scene by manipulation.

Related work demonstrates dense tracking and reconstruction to generate 3D models of static or moving objects [13]–[15]. However, these techniques require a bounding box (or a pre-trained object detector) as system input to define the target.

This paper extends prior work on *unsupervised object discovery, detection, tracking and reconstruction* (DDTR) [11] with a novel form of appearance based object discovery and also uses a robot manipulator to verify candidate objects. A complete dense 3D model is generated once the robot has full observation of an object.

The proposed framework is different from most existing frameworks as: 1) it does not require prior information for object discovery [1], [4]; 2) it builds models and also subsequently detects, tracks and refines *dense* 2D and 3D models [1]–[4], [16], [17] [8]; 3) it learns high quality 3D models of unknown objects via robot-object interaction; 4) it is also a simultaneous localization and mapping system (SLAM), which enables the potential of autonomous object discovery in a large environment; 5) the final learned model of an object can be used for future object detection tasks, and shared with other robots.

The remainder of this paper is structured as follows: Section 2 briefly covers preliminaries and introduces our approach. In the following sections we cover appearance-based object discovery, motion-based objects verification, tracking,

*This work was supported by Google, Toyota and NASA.

Computer Science Department, University of Colorado at Boulder, CO, USA lu.ma|mahsa.ghafarianzadeh|gsibley|david.t.coleman|ncorrell@colorado.edu

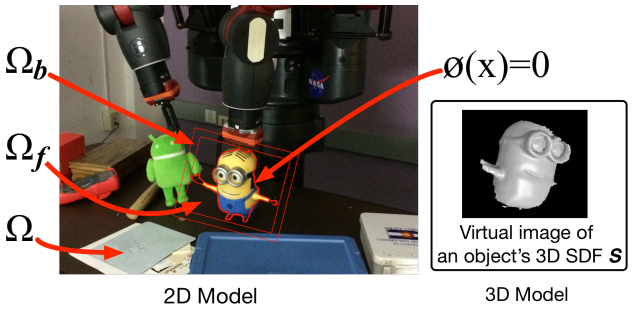

Fig. 2: An Object is represented implicitly with a 2D and a 3D model. The 2D model is represented with a RGB histogram and a 2D SDF $\Phi(x)$. The zero level set of $\Phi(x)$ segments the image domain Ω into a foreground Ω_f and a background Ω_b. The 3D model of the object is represented by the 3D SDF and can be rendered via ray casting and phong shading.

reconstruction and system integration. Section 6 tests and discusses the system performance. Section 7 addresses failure cases and future work and Section 8 draws conclusions.

II. Overview

A. Preliminaries

The proposed system (Fig.1) uses a dual-arm Baxter robot with a calibrated RGBD camera $\mathcal{C}$ mounted on its head, where the world pose of the center of the robot is $T_{wR} \in SE(3)$. The camera $\mathcal{C}$ is calibrated with respect to the robot center, with extrinsic matrix $T_{wC} \in SE(3)$, and intrinsic matrix K_i. The transformation between T_{wR} and T_{wC} is calibrated and considered as a known parameter. An object in the scene is noted as O_i, where its world pose T_{wO} is noted as the center point of the object, and can be computed by $T_{wO} = T_{wC} \oplus T_{CO}$.

A 3D point in the camera frame is denoted as $x_c = (x, y, z)^\top$ and projects onto the image plane as $\pi(x_c) = (x/z, y/z)^\top$ (i.e., dehomogenization). A 2D point $\mathbf{u} = (u, v)^\top$ in the image plane can be back-projected to a 3D point by $\chi = \frac{1}{d} \cdot K_i^{-1} \cdot (\mathbf{u}; 1)$ with a depth value d. The system uses T_{wr}^i and T_{wl}^i to represent the camera pose in the reference frame (time $k-1$) and the live frame (time k).

We denote learned objects as $O = O_\emptyset, O_1, O_2, \cdots, O_n$. Each object is represented by both a 2D and a 3D model. Figure 2 shows the representation of an object's 2D and 3D model.

The system uses shape (2D contour) and appearance (RGB histogram) to represent an object's 2D model. A level-set embedding function $\Phi(x_i)$, namely a 2D signed distance function (SDF), is used to implicitly represent an object's 2D contour (shape). $x_i = x_{i1}, x_{i2} \cdots, x_{in}$ is the set of pixel locations in the coordinate frame of O_i. The zero level-set of the object's 2D SDF $\Phi(x) = 0$ is used to represent the shape (2D contour). The 2D contour $\Phi(x) = 0$ segments the image domain Ω into a foreground Ω_f and a background Ω_b, with appearance models M_f and M_b respectively. Notice that Ω_f^Φ and Ω_b^Φ is used to denote the foreground (pixels inside $\Phi(x) = 0$) and the background segments by the zero level-set of $\Phi(x)$.

The system uses the truncated signed distance function (TSDF) to represent the 3D model $\mathcal{S}_i$ of an object O_i, Where the shape (geometry) of an object can be rendered by ray casting $\mathcal{S}_i$ in different views. The 3D SDF $\mathcal{S}$ is stored as a $n \times n \times n$ volume cube, where n is the number of voxels in one dimension. The size of each voxel is $v_m = \frac{2r}{n}$, where r is the radius of the volume. The 3D SDF is initialized and updated by SDF fusion.

Notice that the 2D contour ($\Phi(x) = 0$) is actually the boundary of the projection of $\mathcal{S}_i$. By using the 2D and 3D models, the system represents objects at different levels and performs 2D&3D tracking and reconstruction Simultaneously during object learning.

B. System Structure

The proposed pipeline goes through the following stages (Fig 3):

Initialization. The system is initialized by creating a 3D model of the first frame (noted as the dominant object $O_\emptyset$) from the input RGBD video. We use a $512 \times 512 \times 512$ volume for $\mathcal{S}_i$.

Appearance-based Object Discovery. For the first $200-300$ frames, the robot moves its head (the camera) to explore the scene which collects enough information for the appearance-based object discovery approach. While the robot explores the scene, the system tracks and updates $O_\emptyset$ frame by frame, and discovers (Section III-A) a set of candidate object contour $C^{'} = C_1^{'}, C_2^{'}, \cdots, C_n^{'}$ by the appearance-based object discovery approach.

Robotic Manipulation. Given a candidate objects contour set $C^{'}$, the system verifies the candidate object $O_i^{'}$ (defined by $C_i^{'} \in C^{'}$) by applying motion (grasping, poking, etc) to $O_i^{'}$. This allow the system generate a verified objects set $O = O_1, O_2 \cdots, O_n$.

Motion-based Object Verification. For a verified object $O_j \in O$, the system generates its contour C_j by extracting portions of the scene that fail to match with $O_\emptyset$ via the ICP+RGB algorithm. The system saves C_j to the discovered objects contour set C.

2D Tracking and 2D Reconstruction. For a discovered object O_j with a contour $C_j \in C$, the system matches C_j with O_i by the appearance, shape and motion cues. This 2D tracking result gives a rough pose estimation of O_j in the image domain of C_j where C_j defines a foreground domain Ω_f^j. 2D reconstruction is then achieved by updating $\Phi(x_i)$ to $\Phi^{'}(x_i)$ in Ω_f^j via LSE.

3D Tracking and 3D Reconstruction. The system uses the ICP+RGB pose estimator for 3D tracking, which estimates the relative camera pose between Ω_f^Φ and $\Omega_f^{\Phi'}$. Once O_i is tracked, the system updates $\mathcal{S}_i$ by fusing every pixel from $\Omega_f^{\Phi'}$ into $\mathcal{S}_i$ via SDF fusion.

III. Unsupervised Object Discovery

The system uses an appearance driven, motion verification pipeline for unsupervised object discovery.

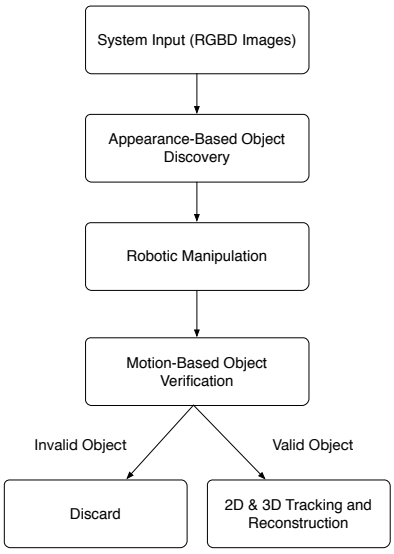

Fig. 3: System flow chart

A. Appearance-Based Object Discovery

The system uses [10] to discover candidate objects in the RGB video sequence. This technique generalizes a graph-based image segmentation to the spatio-temporal case.

In short, given a video sequence, a spatio-temporal graph $G = (V, E, W)$ is constructed, where the graph nodes V are the pixels in the video and are connected by edge E if they are within distance r from each other in each frame, and W measures the similarity of pixels connected by an edge. Also, optical flow is used to add temporal motion information to the graph by connecting pixel (x, y) in frame t to its 9 neighbors along the backward flow (u, v) in frame $t-1$. Affinities are calculated between pixels that are connected in the same frame and also between frames. The overall affinity matrix is a sparse symmetric block diagonal matrix that has the following structure:

$$W = \begin{bmatrix} W_{11} & W_{12} & \dots & 0 \\ W_{21} & W_{22} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & W_{nn} \end{bmatrix} \quad (1)$$

Where W_{ii} is the affinity matrix for frame 1, and W_{ij} is the affinity matrix between frame i and frame j. Then spectral clustering is applied to partition the affinity graph W by computing its k eigenvectors corresponding to the k largest eigenvalues. However to make this process more efficient, only a randomly selected subset of pixels in the video is used to estimate a rank-k approximation of normalized affinity matrix $L = D^{-\frac{1}{2}} W D^{-\frac{1}{2}}$, where D is the diagonal degree matrix defined as $D_{ii} = \sum_{j=1}^{n} W_{ij}$. Next, k-means is used to cluster the eigenvectors which gives an oversegmentation of the video in the form of 3D superpixels. Finally, a post-processing step is applied to merge the 3D superpixels and obtain the final segmentation.

In conclusion, given a RGB image I_k, the appearance-based spatio-temporal object discovery approach $\mathcal{D}(\cdot)$ generates a candidate objects set $C^{'} = C_1^{'}, C_2^{'}, \cdots, C_p^{'}$:

$$C^{'} = \mathcal{D}(I), I = I_1, I_2 \cdots I_K \quad (2)$$

B. Robotic Manipulation

Given a candidate objects set $C^{'}$, the system (robot and the arms) uses a motion cue to verify the validation of an object. The system randomly selects an object and then uses the robot arms to move it (grasping or poking). To interact with the object, the object pose T_{wO} has to be known, which can be computed by:

$$T_{wO} = T_{wC} \oplus T_{CO} \quad (3)$$

where T_{CO} is estimated by the relative pose between the central point of $C_i^{'}$ and the camera. The size of the object can be approximated by estimating the size of $C_i^{'}$ in x, y, z direction. Given the object's pose T_{wO}, a grasp generator determines what grasp and pre-grasp positions are necessary to move the object. An inverse kinematics solver is used to evaluate the feasibility of the generated candidate grasps. If no valid grasp is found due to the geometry or reach-ability of the object, a poking start and end position is generated. A joint trajectory to the pre-grasp or poking start position is generated using the motion planning framework [18]. The trajectory is then executed by the robot. The robot then performs a straight line motion from the pre-grasp position to the grasp position, or from the poking start to poking end position.

If grasping, the end effector constrains the object and lifts it to within proper range of the camera. A pre-defined trajectory of the arm moves (rotates & translates) the object through the scene under the observation of the camera (head mounted). This trajectory is designed to optimize the views of the object that can be observed which allows the system to capture enough data for object learning (tracking and reconstruction).

C. Motion-Based Object Verification

While the robot is manipulating the objects, the system tracks $O_\emptyset$ by estimating the relative pose $T_{rl}^{\emptyset}$ of the camera between the reference frame I_r and the live frame I_l via the ICP+RGB algorithm. A virtual image I_v of $O_\emptyset$ can then be generated by:

$$I_v = \Upsilon(\mathcal{S}_\emptyset, T_{wl}^{\emptyset}) \quad (4)$$

where $\Upsilon(\cdot)$ is the ray casting operator (Section IV-D). If any objects have relative motion against $O_\emptyset$, the system can discover them by extracting portions of the scene that fail to track with $O_\emptyset$:

$$I_o = I_v - I_l \quad (5)$$

By searching the disjoint contours with a minimum number of valid pixels d_1 in I_o, a contour set C of candidate objects can be obtained:

$$C = \Gamma(I_v - I_l) = \Gamma(\Upsilon(\mathcal{S}_\emptyset, T_{wl}^{\emptyset}) - I_l), C = C_1, \cdots, C_p \quad (6)$$

Here $\Gamma(\cdot)$ detects disjoint contours in an image [19].

The motion-based objects verification pipeline allows the system to verify objects from the candidate objects set and discard negative object discovery results if the expected motion between the candidate object and the dominant object does not happen. With this process, the system generates a verified objects set C, which will be used for object learning.

IV. Tracking and Reconstruction

For a discovered (verified) object, the proposed system tracks and reconstructs it in 2D and 3D simultaneously.

A. 2D Tracking

For $C_j \in C$ that matches (tracks) with $O_j \in O$ by the appearance, shape and motion cues [11], C_j defines an image domain Ω^j:

$$\Omega^j = \mathcal{G}(C_j) \tag{7}$$

Here, $\mathcal{G}(C_j)$ inflates the image region defined by C_j. Notice that the 2D tracking process only gives a rough tracking result between C_j and O_j in the 2D image domain.

B. 2D Reconstruction.

Given a discovered object contour $C_i \in C$ (matches with an object O_i), the system does not update $\mathcal{S}_i$ with C_i directly. Instead, it updates the 2D SDF of O_i by level-set-evolution first, and then uses the updated 2D SDF to track and reconstruct O_i in 3D space.

The system avoids updating $\mathcal{S}_i$ with C_i directly because: 1) the precision of the C_i is not guaranteed. Noisy information may be introduced from the object discovery approach to C_i, which would hamper the final 3D reconstruction result if the system updates $\mathcal{S}_i$ with C_i directly. 2) Multiple objects may be overlapping with each other, but it is necessary to extract the precise contours from C_i for each object. 3) The system needs to update the object's 3D model by fusing pixels that have not yet been observed, where C_i only matches against the known object pixels. However, LSE solves this by updating the object's 2D model based on its previous 2D model in C_i.

$$\Phi^{'}(x_i) = \mathcal{L}(\Phi(x_i), \Omega^i) \tag{8}$$

Here $\mathcal{L}(\cdot)$ is the level-set-evolution (LSE) operation [11] that updates the object's 2D SDF $\Phi(x_i)$ to $\Phi(x_i)'$ in $\Omega^i = \mathcal{G}(C_i)$. The foreground (in the RGBD space) defines by $\Phi(x_i)'$ is $\Phi^{'}_f(x_i)$

C. 3D Tracking.

Once the 2D model of an object is updated via LSE, the system tracks $\Phi^{'}_f(x_i)$ against $\Phi_f(x_i)$ for 3D tracking. This is done by estimating the relative camera pose T^i_{rl} between T^i_{wr} (pose of $\Phi_f(x_i)$) and T^i_{wl} (pose of $\Phi^{'}_f(x_i)$) with the ICP+RGB pose estimator $\mathcal{E}(\cdot)$. This combined pose estimator $\mathcal{E}(\cdot)$ allows robust tracking for objects with either complex appearance or shape (geometry).

$$T^i_{rl} = \mathcal{E}(\Phi_f(x_i), \Phi^{'}_f(x_i)) \tag{9}$$

D. 3D Reconstruction.

For an object O_i with the updated 2D SDF $\Phi^{'}(x_i)$, the system updates $\mathcal{S}_i$ by fusing every valid point $\chi = (x, y, z) \in \Phi^{'}_f(x_i)$ into $\mathcal{S}_i$:

$$\mathcal{S}^{'}_i = \mathcal{F}(\mathcal{S}_i, \Phi^{'}_f(x_i), T_{wl}) \tag{10}$$

Here, $\mathcal{F}(\cdot)$ is the SDF Fusion operation. T_{wl} is the global pose of $\Phi^{'}_f(x_i)$. The system also stores intensity in $\mathcal{S}^{'}_i$, where it can render the zero level set surface of $\mathcal{S}$ and generate a virtual gray I^g_v and depth image I^d_v by ray casting $\Upsilon(\cdot)$ [20]:

$$I_v = \Upsilon(\mathcal{S}, T_{wc}), I_v = I^g_v \cup I^d_v \tag{11}$$

where T_{wc} is the pose of the virtual camera.

V. System Integration

Given input RGBD images I, the system discovers a candidate objects set $C^{'}$ in I via the appearance-based spatio-temporal object discovery approach $\mathcal{D}(I)$:

$$C^{'} = \mathcal{D}(I), C^{'} = C^{'}_1, C^{'}_2, \cdots, C^{'}_n \tag{12}$$

For $C^{'}_i \in C^{'}$, the system verifies it by the motion cue (via robot manipulation) and produces a verified objects set C.

$$C = \Gamma(\Upsilon(\mathcal{S}_\emptyset, T^\emptyset_{wl}) - I_l), C = C_1, \cdots, C_n \tag{13}$$

Here, $C_j \in C$ defines an image domain Ω^j that matches with O_i:

$$\Omega^j = \mathcal{G}(C_i) \tag{14}$$

where O_i has camera pose T^i_{wr}, 2D SDF $\Phi(x_i)$ and 3D SDF $\mathcal{S}_i$ in the reference frame. $\Phi(x_i)$ can be updated to $\Phi^{'}(x_i)$ by LSE $\mathcal{L}(\cdot)$.

$$\Phi^{'}(x_i) = \mathcal{L}(\Phi(x_i), \Omega^j) = \mathcal{L}(\Phi(x_i), \mathcal{G}(C_i)) \tag{15}$$

where the foreground defines by $\Phi^{'}(x_i)$ is:

$$\Phi^{'}_f(x_i) = \mathcal{H}(\Phi^{'}(x_i)) = \mathcal{H}(\mathcal{L}(\Phi(x_i), \mathcal{G}(C_i))) \tag{16}$$

here, $\mathcal{H}(\cdot)$ extract pixels inside $\Phi^{'}(x_i) = 0$. Now, 3D tracking is processed by estimating the relative pose T^i_{rl} between $\Phi_f(x_i)$ and $\Phi^{'}_f(x_i)$ via the ICP+RGB pose estimator $\mathcal{E}(\cdot)$:

$$T^i_{rl} = \mathcal{E}(\Phi_f(x_i), \Phi^{'}_f(x_i)) \tag{17}$$

3D Reconstruction is then achieved by updating $\mathcal{S}_i$ with $\Phi^{'}_f(x_i)$ by SDF Fusion $\mathcal{F}(\cdot)$:

$$\begin{aligned}\mathcal{S}'_i = \mathcal{F}(\mathcal{S}_i, \Phi^{'}_f(x_i), T^i_{wl}) = \mathcal{F}(\mathcal{S}_i, \mathcal{H}(\mathcal{L}(\Phi(x_i), \mathcal{G}(C_i))),\\ T^i_{wr} \oplus \mathcal{E}(\Phi_f(x_i), \mathcal{L}(\Phi(x_i), \mathcal{G}(C_i))))\end{aligned} \tag{18}$$

Equation 18 describes how unsupervised object discovery and learning is unified by an appearance driven ($C^{'}$), motion verification (C), tracking (T^i_{wr}) and reconstruction ($\Phi(x_i)$ and $\mathcal{S}_i$) pipeline. Here, the optimization of the whole system is equivalent to the optimization of each subsystem.

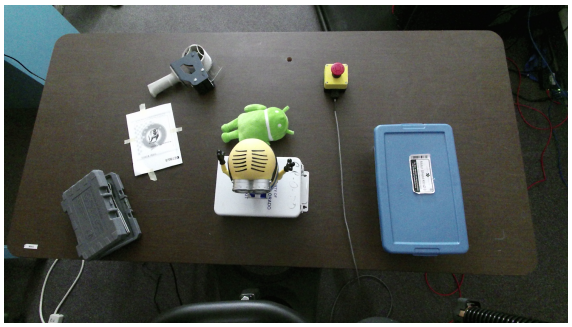

Fig. 4: Test scenario with objects for discovery and modeling by manipulation. The white piece of paper attached to the the desk is considered as a part of the desk (not a separate object) to test the object discovery pipeline.

Fig. 7: Examples of the verified objects set (the five objects under the scene) via poking in the test scenario. The white paper from the appearance based objects discovery result (Fig. 5) is detected as invalid object and discarded. The yellow figure in the middle of the desk is verified and learned by grasping (Fig. 9).

VI. Results

We evaluate the system in an unstructured scenario with different sizes, shapes and colors of objects. Notice that the white piece of paper attached to the middle-left side of the desk is considered as a part of the desk (not a separate object), which is used to test the robustness of the system. The system is tested with RGBD video streams captured at 20 FPS by a Kinect v2 camera. The camera is well calibrated and the RGB and the depth images are aligned during the experiment. Fig. 4 shows an example of the test scenario.

A. Appearance-based object discovery

We compare our appearance-based object discovery method with two state-of-the-art object discovery approaches [21], [22]. Fig. 5 shows an example of the final discovered objects. As Fig. 5 shows, [21] is able to generate good object discovery results (Fig. 5(k)) but its performance is not stable and robust (Fig. 5(h), 5(l)). Meanwhile, [22] generates stable object discovery results but the quality is not guaranteed (fail to discover the gray box in all sequences). The proposed method instead generates stable and precise object discovery results, which gives a good hint for the motion-based object verification pipeline.

Appearance-based object discovery methods are usually used for image segmentation, which are not designed for indoor scene object discovery. To evaluate such methods, we use a pixel level per-frame precision and recall curve approach. Here, the ground truth 3D model $\mathcal{S}_i$ of an object O_i is approximated once the system learned a complete 3D model of the object, as the motion-based object discovery and learning pipeline is very robust and precise [11]. In each subsequent frames, the system uses the virtual image (D_v) of $\mathcal{S}_i$ to approximate the ground truth label and uses k different appearance-based object discovery methods to generate different discovered objects images D_l^k. The performance of the system is evaluated at pixel level by comparing each corresponding object's pixel p in D_v and D_l^k. The system considers a good match (true positive) if D_v and D_l align (match).

We evaluate appearance-based object discovery methods with a poking and a grasping example and report the per-frame precision and recall curve (Fig. 6). As Fig. 6 shows, in general our appearance-based object discovery method (the red curve) has the best performance.

B. Motion-Based Objects Verification

Taking the appearance-based object discovery results as a hint, the robot verifies each candidate object by motion. We test the system with a poking (the gray box, the white paper, the tape dispenser, the yellow button, the green Android and the blue box) and a grasping (the yellow cartoon figure) experiment. Fig. 7 shows an example of the final verified objects in the poking experiment. (a grasping example is available at Fig. 9).

Due to the limitation of the appearance-based object discovery approach, negative results (e.g. the white paper on the desk) will be produced as it considers textures (color) to be candidate objects. However, such negative discovery results can be correctly recognized by the motion-based verification pipeline. The system will discard negative objects discovery result and only focus on the positive object discovery results. Notice that the motion-based objects verification pipeline does not discovery the blue box completely as its initial position (Fig. 4) and final position overlaps (with the same geometry and appearance). However, the system is able recovery it via the tracking and reconstruction approaches and learn its full 3D model (Fig. 8).

C. Learning

Once an object is verified, the system learns the 2D and 3D model of the object via the tracking and reconstruction pipeline. As Fig. 8 shows, the system is capable to learn all candidate objects verified by the poking experiment.

To obtain a complete model of an object, the observation of multi-views of the object is required. We demonstrate this by a grasping experiment, where the robot grasps an object (the yellow cartoon figure) on the desk and moves it

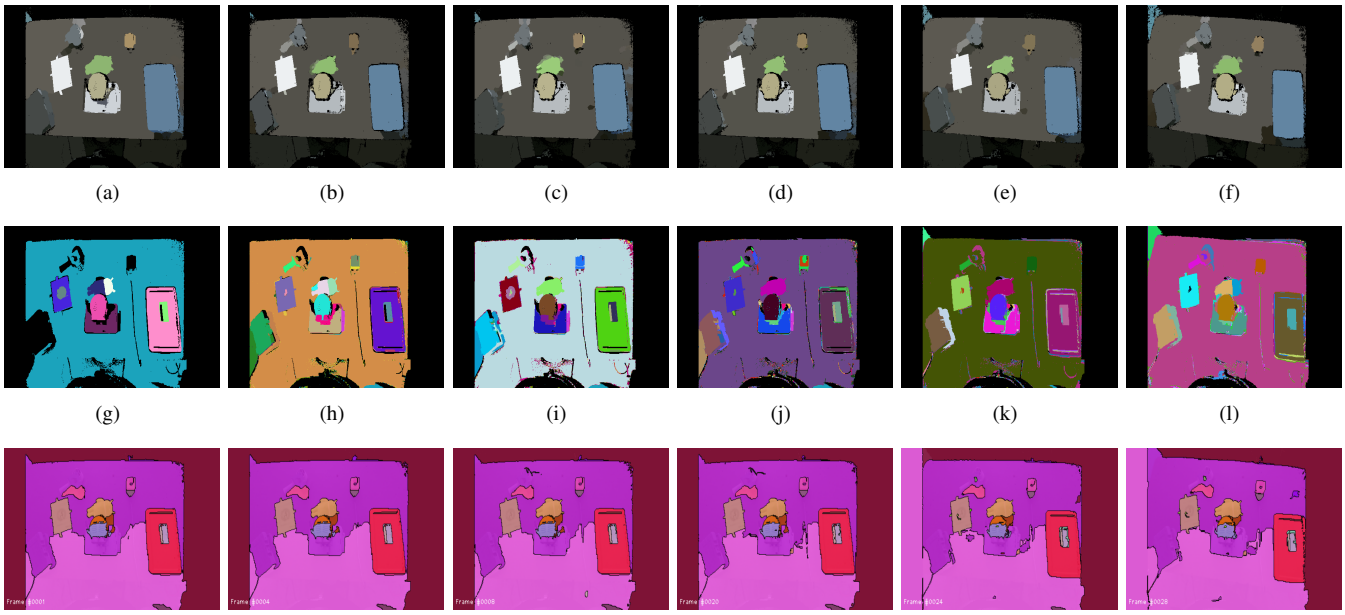

Fig. 5: Appearance-based object discovery results by the proposed method (the first row), [21] (the second row) and [22] (the third row) . The proposed method has the best performance in terms of precision and recall (Fig 6) of object pixels (scored using motion-verified object boundaries).

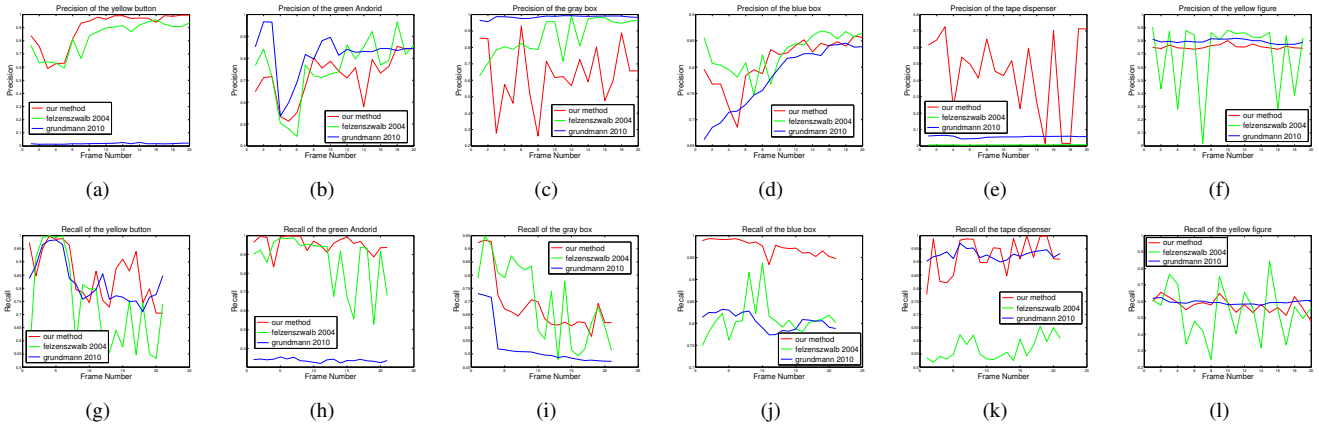

Fig. 6: Pixel level per-frame precision (the first row) and recall (the second row) curve of our approach (the red curve), [21] (the green curve) and [22] (the blue curve). In general, our approach has the best performance in the test scenario.

under the observation of the camera. A complete 3D model of the object is reconstructed when the system has a complete knowledge of the object (takes approximate 500 frames). Fig. 9 shows an example of the final learning result of the object via grasping.

D. Time Analysis

The system is tested using a single NVIDIA TITAN GPU, Intel i7 CPU desktop, with 640×480 resolution of input images. Table 1 shows our system run-time in different stages. Notice that the robotic manipulation, motion-based object verification, tracking and reconstruction approaches are processed simultaneously. The proposed system has a low running time complexity (22 Hz/object) during the learning (tracking and reconstruction) stages (GPU-based). The appearance-based object discovery approach is relatively slow, which takes approximately 10 min to discover objects using 10 frames (CPU-based). However, a similar CPU-based system [9] has been implemented which runs at 1 Hz. Based on the time analysis above, it is reasonable to conjecture that our system can run in real-time after optimization.

VII. Failure Cases and Future Work

Although the system is robust to many real-world operating conditions, there are several limitations of our current work. The appearance-based object discovery method

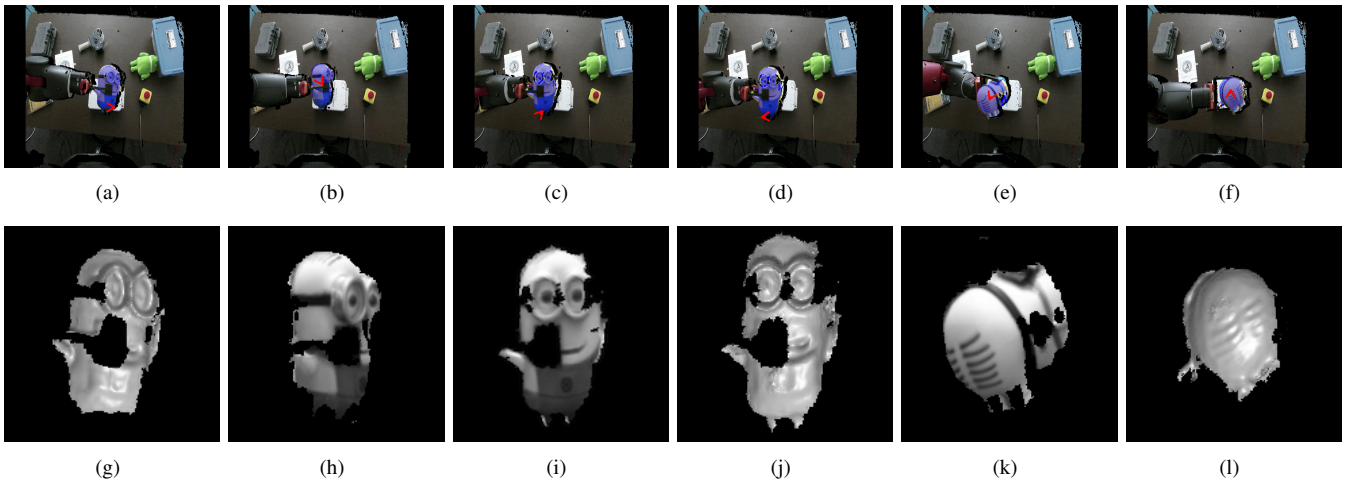

Fig. 9: Tracking (the first row, shown by blue transparent masks) and reconstruction (the second row, shown by the virtual image of the object's 3D SDF) results of a verified object (the yellow cartoon figure). In this scenario, the robot grasps an object on the desk and moves it around the scene (for 500 frames). A complete 3D model of the object is learned after the robot has enough observation of the object. The red arrow shows the object's estimated velocities.

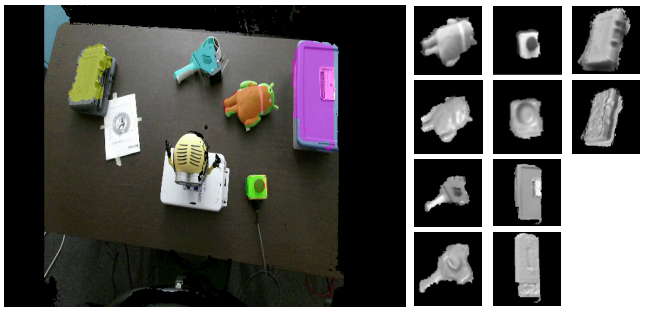

Fig. 8: Final learned objects of the poking experiment. The left figure shows the final tracking results (shown as different color transparent masks). The right figure shows the virtual gray and the phong shading images of the final reconstruction results.

TABLE I: System run-time

Appearance-Based Object Discovery (10 frames)	10 min
Manipulation Path Planning	5 s
Robotic Manipulation	1 min
Motion-Based Object Verification	10 ms
Tracking & Reconstruction (one object)	38 ms

depends on the difference between the foreground and the background and cannot discover objects under extremely dark/bright environments. For robot manipulation and motion-based object verification approaches, due to the limitation of the robot platform we have, the system cannot grasp or poke too small, too large, or too heavy objects. Additionally, the learning pipeline requires sufficient pixels for robust tracking, which means that sometimes the system will not learn the model for very small objects.

The success of our framework opens many areas of future work. The appearance-based object discovery pipeline can be improved by integrating object shape information. The system could use a single RGB camera by using techniques from [23], [24]. The path planning for the grasping/learning stage can be improved by a feedback system between the vision and the manipulation pipeline. The learned 3D model of objects could be shared among different robots for similar object detection by 3D tracking techniques [25].

VIII. Conclusion

We present a novel unsupervised framework for a robot to discover and learn unknown objects in the scene by manipulation. The system performs dense 3D simultaneous localization and mapping concurrently with unsupervised object discovery. Spatio-temporal cues and appearance cues are used to produce a set of candidate objects. A motion-based verification strategy is used to verify candidate objects, and object appearance and shape is subsequently learned by grasping and poking. We compare three different approaches for appearance-based object discovery and find that a novel form of spatio-temporal super-pixels gives the highest quality candidate object models in terms of precision and recall.

References

[1] L. Bo, X. Ren, and D. Fox, "Unsupervised feature learning for rgb-d based object recognition," in *Experimental Robotics*. Springer, 2013, pp. 387–402.

[2] F. Endres, C. Plagemann, C. Stachniss, and W. Burgard, "Unsupervised discovery of object classes from range data using latent dirichlet allocation." in *Robotics: Science and Systems*, vol. 2. Seattle, Washington;, 2009.

[3] E. Herbst, X. Ren, and D. Fox, "Rgb-d object discovery via multi-scene analysis," in *Intelligent Robots and Systems (IROS), IEEE/RSJ Int. Conf. on*, 2011, pp. 4850–4856.

[4] D. Meger, P.-E. Forssén, K. Lai, S. Helmer, S. McCann, T. Southey, M. Baumann, J. J. Little, and D. G. Lowe, "Curious george: An attentive semantic robot," *Robotics and Autonomous Systems*, vol. 56, no. 6, pp. 503–511, 2008.

[5] B. Morisset, R. B. Rusu, A. Sundaresan, K. Hauser, M. Agrawal, J.-C. Latombe, and M. Beetz, "Leaving flatland: toward real-time 3d navigation," in *Robotics and Automation (ICRA), IEEE Int. Conf. on*, 2009, pp. 3786–3793.

[6] K. Lai, L. Bo, X. Ren, and D. Fox, "A large-scale hierarchical multi-view rgb-d object dataset," in *Robotics and Automation (ICRA), IEEE Int. Conf. on*, 2011, pp. 1817–1824.

[7] E. Herbst, P. Henry, X. Ren, and D. Fox, "Toward object discovery and modeling via 3-d scene comparison," in *Robotics and Automation (ICRA), 2011 IEEE Int. Conf. on*, 2011, pp. 2623–2629.

[8] A. Collet, B. Xiong, C. Gurau, M. Hebert, and S. S. Srinivasa, "Exploiting domain knowledge for object discovery," in *Robotics and Automation (ICRA), 2013 IEEE Int. Conf. on*, 2013, pp. 2118–2125.

[9] M. Grundmann, V. Kwatra, M. Han, and I. Essa, "Efficient hierarchical graph-based video segmentation," in *Computer Vision and Pattern Recognition (CVPR), IEEE Conference on*, 2010, pp. 2141–2148.

[10] M. Ghafarianzadeh, M. Blaschko, G. Sibley, *et al.*, "Unsupervised spatio-temporal segmentation with sparse spectral clustering," in *British Machine Vision Conference (BMVC)*, 2014.

[11] L. Ma and G. Sibley, "Unsupervised dense object discovery, detection, tracking and reconstruction," in *Computer Vision–ECCV 2014*, 2014, pp. 80–95.

[12] J. Mason, B. Marthi, and R. Parr, "Object disappearance for object discovery," in *Intelligent Robots and Systems (IROS), IEEE/RSJ Int. Conf. on*, 2012, pp. 2836–2843.

[13] C. Y. Ren, V. Prisacariu, D. Murray, and I. Reid, "STAR3D: Simultaneous tracking and reconstruction of 3D objects using RGB-D data."

[14] A. Dame, V. A. Prisacariu, C. Y. Ren, and I. Reid, "Dense reconstruction using 3d object shape priors," in *Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on*. IEEE, 2013, pp. 1288–1295.

[15] S. Y. Bao, M. Chandraker, Y. Lin, and S. Savarese, "Dense object reconstruction with semantic priors," in *Computer Vision and Pattern Recognition (CVPR), IEEE Conference on*, 2013, pp. 1264–1271.

[16] R. Triebel, J. Shin, R. Siegwart, R. Y. Siegwart, and R. Y. Siegwart, "Segmentation and unsupervised part-based discovery of repetitive objects," in *Robotics: Science and Systems*, vol. 2, 2010.

[17] J. Modayil and B. Kuipers, "Bootstrap learning for object discovery," in *Intelligent Robots and Systems, 2004.(IROS 2004). Proceedings. 2004 IEEE/RSJ Int. Conf. on*, vol. 1, 2004, pp. 742–747.

[18] D. Coleman, I. Sucan, S. Chitta, and N. Correll, "Reducing the barrier to entry of complex robotic software: a moveit! case study," *Journal of Software Engineering for Robotics*, 2014.

[19] C.-H. Teh and R. T. Chin, "On the detection of dominant points on digital curves," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 11, no. 8, pp. 859–872, 1989.

[20] R. A. Newcombe, A. J. Davison, S. Izadi, P. Kohli, O. Hilliges, J. Shotton, D. Molyneaux, S. Hodges, D. Kim, and A. Fitzgibbon, "Kinectfusion: Real-time dense surface mapping and tracking," in *Mixed and augmented reality (ISMAR), 10th IEEE Int. Symp. on*, 2011, pp. 127–136.

[21] P. F. Felzenszwalb and D. P. Huttenlocher, "Efficient graph-based image segmentation," *International Journal of Computer Vision*, vol. 59, no. 2, pp. 167–181, 2004.

[22] M. Grundmann, V. Kwatra, M. Han, and I. Essa, "Efficient hierarchical graph based video segmentation," *IEEE CVPR*, 2010.

[23] R. A. Newcombe, S. J. Lovegrove, and A. J. Davison, "Dtam: Dense tracking and mapping in real-time," in *Computer Vision (ICCV), IEEE Int. Conf. on*, 2011, pp. 2320–2327.

[24] R. A. Newcombe and A. J. Davison, "Live dense reconstruction with a single moving camera," in *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, 2010, pp. 1498–1505.

[25] V. A. Prisacariu and I. D. Reid, "Pwp3d: Real-time segmentation and tracking of 3d objects," *International journal of computer vision*, vol. 98, no. 3, pp. 335–354, 2012.